Sparse Recovery with Graph Constraints

Meng Wang, Weiyu Xu, Enrique Mallada, Ao Tang

arXiv:1207.2829v1 [cs.IT] 12 Jul 2012

***Abstract*—Motivated by the need to monitor large-scale networks, this paper addresses the problem of sparse recovery with graph constraints. Unlike conventional sparse recovery where a measurement can contain any subset of the unknown variables(nodes), we take an additive measurement over nodes only if they induce a connected subgraph. We provide explicit measurement constructions for several special graphs, and the number of measurements by our construction is less than that needed by existing random constructions. Moreover, our construction for a line network is provably optimal in the sense that it requires the minimum number of measurements. A measurement construction algorithm for general graphs is also proposed and evaluated. For any given graph G with n nodes, we derive bounds of the minimum number of measurements needed to recover any k-sparse vector over G ($M_{k,n}^G$). Using the Erdős-Rényi random graph as an example, we characterize the dependence of $M_{k,n}^G$ on the graph structure. Our study suggests that $M_{k,n}^G$ may serve as a graph connectivity metric.**

***Index Terms*—sparse recovery, compressed sensing, topological graph constraints, measurement construction.**

I. Introduction

Network monitoring is a critical module in the operation and management of communication networks, where one keeps track of network state parameters, such as bandwidth utilization and queueing delay. Since measuring each component (e.g., router) in the network directly can be operationally costly, if feasible at all, the topic of inferring system internal characteristics from indirect end-to-end (aggregate) measurements becomes important. This area is known as network tomography, and has been extensively studied during the last decade or so [9], [12], [15], [23], [27], [31], [39].

In many cases, the total number of aggregate measurements is much smaller than the number of components in a network. But we still hope to extract the status of each individual component with some prior knowledge of the unknown signal to recover. For instance, if the signal is sparse, i.e. most entries are zero, we can recover it exactly even though the number of measurements is much smaller than its dimension. For example, transmission delays in the communication networks can be represented by an approximately sparse signal, since only a small number of bottleneck links experience large delays, while the delay is approximately zero elsewhere. *Sparse Recovery* addresses the problem of recovering sparse signals from a smaller number of measurements, and has two different but closely related problem formulations. One is *Compressed Sensing* [5], [10], [11], [19], [20], [25], where the signal is represented by a high-dimensional real vector, and an aggregate measurement is the arithmetical sum of the corresponding real entries. The other is *Group Testing* [21], [22], where the high-dimensional signal is binary and a measurement is a logical disjunction (**OR**) on the corresponding binary values.

One key question in both compressed sensing and group testing is to design a small number of non-adaptive measurements (either real or logical) such that all the vectors (either real or logical) up to certain sparsity (the support size of a vector) can be correctly recovered. Most existing results, however, rely critically on the assumption that any subset of the values can be aggregated together [10], [19], which is not realistic in network monitoring problems where only objects that form a path or a cycle on the graph [1], [27], or induce a connected subgraph can be aggregated together in the same measurement. Only a few recent works consider graph topological constraints, either in group testing [13] setup, especially motivated by link failure localization in all-optimal networks [3], [13], [28], [32], [36], or in compressed sensing setup, with application in estimation of network parameters [16], [25], [29], [38].

Though motivated by network monitoring problems, sparse recovery with graph constraints abstractly models scenarios when certain elements cannot be measured together in a complex system. These constraints can result from various reasons, not necessarily lack of connectivity. Therefore, our results can be potentially useful to other applications besides network tomography.

Here are the main contributions of this paper.

(1) We provide explicit measurement constructions for various graphs. Our construction for line networks is optimal in the sense that it requires the minimum number of measurements. For other special graphs, the number measurements by our construction is less than the existing estimates (e.g. [13], [38]) of the number of measurements required to recover sparse vectors over graphs. (Section III)

(2) We propose a measurement design guideline based on r*-partition* for general graphs and further show some of its properties. (Section IV-A)

(3) A simple measurement design algorithm is proposed for general graphs, and we evaluate its performance both theoretically and numerically. (Section IV-B and VIII)

(4) For Erdős-Rényi random graphs, we characterize the dependence of the number of measurements for sparse recovery on the graph structure. (Section V)

(5) Motivated by practical needs, we further propose measurement construction methods under additional graph constraints including measurement length constraint, and the requirement that each measurement should pass one of a fixed set of nodes. (Section VI) We also address the issue of sparse recovery when some critical measurements may contain errors. (Section VII)

M. Wang, E. Mallada and A. Tang are with Cornell University, Ithaca, NY. W. Xu is with the University of Iowa, Iowa City, IA.

Partial and preliminary results have appeared in [34].

II. Model and Problem Formulation

Consider a graph $G = (V, E)$, where V denotes the set of nodes with cardinality $|V| = n$ and E denotes the set of links. Each node i is associated with a real number x_i, and we say vector $\mathbf{x} = (x_i, i = 1, ..., n)$ is associated with G. Let $T = \{i \mid x_i \neq 0\}$ denote the support of $\mathbf{x}$, and let $\|\mathbf{x}\|_0 = |T|$ denote the number of non-zero entries of $\mathbf{x}$, we say $\mathbf{x}$ is a k-sparse vector if $\|\mathbf{x}\|_0 = k$.

Let $S \subseteq V$ denote a subset of nodes in G. Let E_S denote the subset of links with both ends in S, then $G_S = (S, E_S)$ is the induced subgraph of G. We have the following two assumptions throughout the paper:

(A1): A set S of nodes can be measured together in one measurement if and only if G_S is connected.

(A2): The measurement is an additive sum of values at the corresponding nodes.

(A1) captures the graph constraints. One practical example is a sensor network where the nodes represent sensors and the links represent feasible communication between sensors. For a set S of nodes that induce a connected subgraph, one node u in S, we call u an "agent", monitors the sum of node values in S. Every node in S obtains values from its children, if any, on the spanning tree rooted at u, sums them up with its own value and sends the sum to its parent. Then the central operator can obtain the sum of node values in S by only communicating with the agent u. (A2) follows from the additive property of many network characteristics, e.g. delays and packet loss rates [27]. Note that compressed sensing can also be applied to cases where (A2) does not hold, e.g., the measurements can be nonlinear as in [6], [33].

Let $\mathbf{y} \in \mathcal{R}^m$ $(m \ll n)$ denote the vector of m measurements. A is an $m \times n$ measurement matrix with its ith row corresponding to the ith measurement, i.e., $A_{ij} = 1$ $(i = 1, ..., m, j = 1, ..., n)$ if and only if node j is included in the ith measurement and $A_{ij} = 0$ otherwise. We can write in the compact form that $\mathbf{y} = A\mathbf{x}$. We say a measurement matrix A can ***identify all*** k***-sparse vectors*** if and only if $A\mathbf{x}_1 \neq A\mathbf{x}_2$ for every two different vectors $\mathbf{x}_1$ and $\mathbf{x}_2$ that are at most k-sparse. This definition indicates that every vector $\mathbf{x}$ that is at most k-sparse can be recovered from $A\mathbf{x}$ via ℓ_0-minimization, which returns the sparsest vector among all vectors that can produce the same observation $A\mathbf{x}$. Sparse recovery theory indicates that it is possible to identify n-dimensional vectors from m $(m \ll n)$ measurements provided that the vectors are sparse enough.

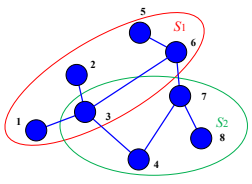


Fig. 1. Network Example

With the above assumptions, A is a 0-1 matrix, and for each row of A, the set of nodes that correspond to '1' should form a connected induced subgraph of G. For example in Fig. 1, we can measure the sum of nodes in S_1 and S_2 by two separate measurements, and the measurement matrix is

$$A = \begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}.$$

We remark here that in group testing with graph constraints, the requirements for the measurement matrix A are the same, while group testing differs from compressed sensing only in that (1) $\mathbf{x}$ is a logical vector, and (2) the operations used in each group testing measurement are the logical "AND" and "OR". Here we consider compressed sensing if not otherwise specified, and the main results are stated in theorems. We sometimes discuss group testing for comparison, and the results are stated in propositions. Note that for recovering 1-sparse vectors, the numbers of measurements required by compressed sensing and group testing are the same.

Given a graph G with n nodes, let $M_{k,n}^G$ denote the minimum number of non-adaptive measurements needed to identify all k-sparse vectors associated with G. Let $M_{k,n}^C$ denote the minimum number of non-adaptive measurements needed in a complete graph with n nodes. Since in a complete graph, any subset of nodes forms a connected subgraph, every 0-1 matrix is a feasible measurement matrix there. Existing results [5], [11], [37] show that with overwhelming probability a random 0-1 matrix with $O(k \log(n/k))$ rows[1] can identify all k-sparse vectors associated with a complete graph, and we can recover the sparse vector by ℓ_1-minimization, which returns the vector with the least ℓ_1-norm[2] among those that can produce the obtained measurements. Then we have

$$M_{k,n}^C = O(k \log(n/k)). \tag{1}$$

Explicit constructions of measurement matrices for complete graphs also exist, e.g., [2], [5], [17], [18], [37]. We use $f(k, n)$ to denote the number of measurements to recover k-sparse vectors associated with a complete graph of n nodes by a particular measurement construction method. $f(k, n)$ varies for different construction methods, and clearly $f(k, n) \geq M_{k,n}^C$. Table I summarizes the key notations.

The questions we would like to address in the paper are:

- Given graph G, what is the corresponding $M_{k,n}^G$? What is the dependence of $M_{k,n}^G$ on G?

TABLE I
SUMMARY OF KEY NOTATIONS

Notation	Meaning
G_S	Subgraph of G induced by S
$M_{k,n}^G$	Minimum number of measurements needed to identify k-sparse vectors associated with G of n nodes.
$M_{k,n}^C$	Minimum number of measurements needed to identify k-sparse vectors associated with a complete graph of n nodes.
$f(k, n)$	Number of measurements constructed to identify k-sparse vectors associated with a complete graph of n nodes

[1]We use the notations $g(n) \in O(h(n))$, $g(n) \in \Omega(h(n))$, or $g(n) = \Theta(h(n))$ if as n goes to infinity, $g(n) \leq ch(n)$, $g(n) \geq ch(n)$ or $c_1 h(n) \leq g(n) \leq c_2 h(n)$ eventually holds for some positive constants c, c_1 and c_2 respectively.

[2]The ℓ_p-norm $(p \geq 1)$ of $\mathbf{x}$ is $\|\mathbf{x}\|_p = (\sum_i |x_i|^p)^{1/p}$, $\|\mathbf{x}\|_\infty = \max_i |x_i|$, and $\|\mathbf{x}\|_0 = |\{i : x_i \neq 0\}|$.

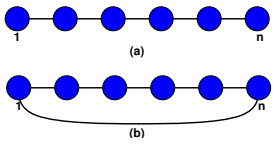


Fig. 2. (a) line network (b) ring network

- How can we explicitly design measurements such that the total number of measurements is close to $M_{k,n}^G$?

III. Sparse Recovery over Special Graphs

In this section, we consider four kinds of special graphs: one-dimensional line/ring network, ring with each node connecting to its four closest neighbors, two-dimensional grid and a tree. The measurement construction method for a line/ring network is different from those for the other graphs, and our construction is optimal (or near optimal) for a line (or ring) network. For other special graphs, we construct measurements based on the "hub" idea and will later extend it to general graphs in Section IV.

A. Line and Ring

First consider a line/ring network as shown in Fig. 2. When later comparing the results here with those in Section III-B, one can see that the number of measurements required for sparse recovery can be significantly different in two graphs that only differ from each other with a small number of links.

In a line/ring network, there is not much freedom in the measurement design since only consecutive nodes can be measured together from (A1). Recovering 1-sparse vectors associated with a line (or ring) network with n nodes is considered in [28], [32], which show that $\lceil \frac{n+1}{2} \rceil$ (or $\lceil \frac{n}{2} \rceil$) measurements are both necessary and sufficient in this case. Here, we consider recovering k-sparse vectors for $k \geq 2$.

Our construction works as follows. Given k and n, let $t = \lfloor \frac{n+1}{k+1} \rfloor$. We construct $n+1-\lfloor \frac{n+1}{k+1} \rfloor$ measurements with the ith measurement passing all the nodes from i to $i+t-1$. Let $A^{(n+1-t)\times n}$ be the corresponding measurement matrix, then its ith row has '1's from entry i to entry $i+t-1$ and '0's elsewhere. For example, when $k=3$ and $n=11$, we have $t=3$, and

$$A = \begin{bmatrix} 1&1&1&0&0&0&0&0&0&0&0 \\ 0&1&1&1&0&0&0&0&0&0&0 \\ 0&0&1&1&1&0&0&0&0&0&0 \\ 0&0&0&1&1&1&0&0&0&0&0 \\ 0&0&0&0&1&1&1&0&0&0&0 \\ 0&0&0&0&0&1&1&1&0&0&0 \\ 0&0&0&0&0&0&1&1&1&0&0 \\ 0&0&0&0&0&0&0&1&1&1&0 \\ 0&0&0&0&0&0&0&0&1&1&1 \end{bmatrix}. \quad (2)$$

Let $M_{k,n}^L$ and $M_{k,n}^R$ denote the minimum number of measurements required to recover k-sparse vectors in a line/ring network respectively. We have

Theorem 1. *Our constructed $n+1-\lfloor \frac{n+1}{k+1} \rfloor$ measurements can identify all k-sparse vectors associated with a line/ring network with n nodes. Moreover,*

$$M_{k,n}^L = n+1-\lfloor \frac{n+1}{k+1} \rfloor \leq M_{k,n}^R + 1. \quad (3)$$

Proof: We first prove that the constructed $n+1-\lfloor \frac{n+1}{k+1} \rfloor$ measurements can identify all k-sparse vectors associated with a line/ring network with n nodes.

The recovery is successful if and only if every $2k$ columns of A are linearly independent, i.e., every non-zero vector $\mathbf{z}$ such that $A\mathbf{z} = \mathbf{0}$ has at least $2k+1$ non-zero elements [10].

A is an identity matrix when $t=1$, and the argument holds trivially. We focus on the case that $t \geq 2$. Now $tk \leq n-1$ holds from definition. For each index $1 \leq k' \leq k-1$, define a submatrix $A_{k'}$, which consists of the first $tk'+1$ rows and the first $tk'+t$ columns of A. For example, for A in (2),

$$A_1 = \begin{bmatrix} 1&1&1&0&0&0 \\ 0&1&1&1&0&0 \\ 0&0&1&1&1&0 \\ 0&0&0&1&1&1 \end{bmatrix},$$

$$A_2 = \begin{bmatrix} 1&1&1&0&0&0&0&0&0 \\ 0&1&1&1&0&0&0&0&0 \\ 0&0&1&1&1&0&0&0&0 \\ 0&0&0&1&1&1&0&0&0 \\ 0&0&0&0&1&1&1&0&0 \\ 0&0&0&0&0&1&1&1&0 \\ 0&0&0&0&0&0&1&1&1 \end{bmatrix}.$$

We claim that: (*) every non-zero vector $\mathbf{w}$ such that $A_{k'}\mathbf{w} = \mathbf{0}$ has at least $2k'+2$ non-zero elements, and at least two non-zero elements exist in its last t entries. We will prove (*) by induction over k'.

First consider A_1. Let $\mathbf{a}_i^T$ denote its ith row. Because any two columns of A_1 are linearly independent, any $\mathbf{w} \neq \mathbf{0}$ such that $A_1\mathbf{w} = \mathbf{0}$ must have at least three non-zero elements. Let j be the index of the last non-zero element of $\mathbf{w}$. Suppose $j \leq t$, then $\mathbf{a}_j^T\mathbf{w} = w_j \neq 0$, contradicting the fact that $A_1\mathbf{w} = \mathbf{0}$. Then $j \geq t+1$ must hold, i.e., at least one of the last t entries of $\mathbf{w}$ is non-zero. From $\mathbf{a}_{t+1}^T\mathbf{w} = 0$, at least two non-zero elements exist in the last t entries of $\mathbf{w}$. One can similarly argue that at least two non-zero elements exist in the first t entries of $\mathbf{w}$. Thus, (*) holds for A_1.

Now suppose (*) holds for $A_{k'}$ with integer k' in $[1, k-2]$. Consider $\mathbf{w} \neq \mathbf{0}$ such that $A_{k'+1}\mathbf{w} = \mathbf{0}$. Let $\hat{\mathbf{w}}$ denote the subvector of the first $tk'+t$ entries of $\mathbf{w}$, then $A_{k'}\hat{\mathbf{w}} = \mathbf{0}$. We remark that $\hat{\mathbf{w}} \neq \mathbf{0}$. Suppose $\hat{\mathbf{w}} = \mathbf{0}$, and let j denote the index of the first non-zero entry of $\mathbf{w}$. Then the inner prodcuct of the $(j+1-t)$th row of $A_{k'+1}$ with $\mathbf{w}$ equals to w_j, which is non-zero, contracting the fact that $A_{k'+1}\mathbf{w} = \mathbf{0}$.

Since $\hat{\mathbf{w}} \neq \mathbf{0}$, from the induction hypothesis, it has at least $2k'+2$ non-zero elements, and at least two non-zero elements in its last t entries. Now consider the last $2t$ entries of $\mathbf{w}$ and the last $t+1$ rows of $A_{k'+1}$. From a similar argument as for A_1, we know that $\mathbf{w}$ must have at least two non-zero elements in the last t entries. So $\mathbf{w}$ has at least $2(k'+1)+2$ non-zero elements, and (*) holds for $k'+1$.

By induction over k', we conclude that (*) follows. Consider any $\mathbf{z} \neq \mathbf{0}$ such that $A\mathbf{z} = \mathbf{0}$. Let $\mathbf{w}$ contain its first kt entries. One can argue that $\mathbf{w} \neq \mathbf{0}$. Then by (*), $\mathbf{w}$ has at least $2k$ non-zero entries. Let j denote the index of the last non-zero entry of $\mathbf{w}$. Since the inner product between $\mathbf{z}$ and the jth row of A is zero, there exists at least one non-zero entry in the last $n-kt$ entries of $\mathbf{z}$. Thus, $\mathbf{z}$ has at least $2k+1$ non-zero entries in total. This concludes the proof.

We next prove that the number of measurements needed to recover k-sparse vectors associated with a line (or ring) networks is at least $n+1-\lfloor\frac{n+1}{k+1}\rfloor$ (or $n-\lfloor\frac{n}{k+1}\rfloor$.)

Let $A^{m\times n}$ denote a measurement matrix with which one can recover k-sparse vectors associated with a line network with n nodes. Then every $2k$ columns of A must be linearly independent. We will prove that $m \geq n+1-\lfloor\frac{n+1}{k+1}\rfloor$.

Let $\boldsymbol{\beta}^i$ denote the ith column of A. Define $\boldsymbol{\alpha}^1 = \boldsymbol{\beta}^1$, $\boldsymbol{\alpha}^i = \boldsymbol{\beta}^i - \boldsymbol{\beta}^{i-1}$ for all $2 \leq i \leq n$, and $\boldsymbol{\alpha}^{n+1} = -\boldsymbol{\beta}^n$. Define matrix $P^{m\times(n+1)} = (\boldsymbol{\alpha}^i, 1 \leq i \leq n+1)$. Since A is a measurement matrix for a line network, each row of P contains one '$1'$ entry and one '$-1'$ entry, and all the other entries must be '0's.

Given P, we construct a graph G_{eq} with $n+1$ nodes as follows. For every row i of P, there is an edge (j,k) in G_{eq}, where $P_{ij} = 1$ and $P_{ik} = -1$. Then G_{eq} contains m edges, and P can be viewed as the transpose of an oriented incidence matrix of G_{eq}. Let S denote the set of indices of nodes in a component of G_{eq}, then one can check that

$$\sum_{i\in S} \boldsymbol{\beta}^i = \mathbf{0}. \tag{4}$$

Since every $2k$ columns of A are linearly independent, every k columns of P are linearly independent, which then implies that the sum of any k columns of P is not a zero vector. With (4), we know that any component of G_{eq} should have at least $k+1$ nodes. Since a component with r nodes contains at least $r-1$ edges, and G_{eq} has at most $\lfloor\frac{n+1}{k+1}\rfloor$ components, then G_{eq} contains at least $n+1-\lfloor\frac{n+1}{k+1}\rfloor$ edges. The claim follows.

We next consider the ring network. Let $\tilde{A}$ denote the measurement matrix with which one can recover k-sparse vectors on a ring network with n nodes. Let $\tilde{\beta}^i$ denote the ith column of $\tilde{A}$. Define $\tilde{\alpha}^1 = \tilde{\beta}^1 - \tilde{\beta}^n$, and $\tilde{\alpha}^i = \tilde{\beta}^i - \tilde{\beta}^{i-1}$ for all $2 \leq i \leq n$. Define matrix $\tilde{P}^{m\times n} = (\tilde{\alpha}^i, 1 \leq i \leq n)$. Similarly, we construct a graph $\tilde{G}_{eq}$ with n nodes based on $\tilde{P}$, and each component of $\tilde{G}_{eq}$ should have at least $k+1$ nodes. Thus, $\tilde{G}_{eq}$ contains at most $\lfloor\frac{n}{k+1}\rfloor$ components, and has at least $n-\lfloor\frac{n}{k+1}\rfloor$ edges. Then we have

$$M^R_{k,n} \geq n - \lfloor\frac{n}{k+1}\rfloor \geq n - \lfloor\frac{n+1}{k+1}\rfloor,$$

and the inequality of (3) holds. ■

Theorem 1 indicates that our construction is optimal for a line network in the sense that the number of measurements is equal to the minimum needed to recover k-sparse vectors. It is near optimal for a ring network, since the number of measurements is no more than the minimum plus one. This improves over our earlier result (Theorem 1 in [34]), which does not have optimality guarantee.

We can save about $\lfloor\frac{n+1}{k+1}\rfloor - 1$ measurements but still be able to recover k-sparse vectors in a line/ring network via compressed sensing. But for group testing on a line/ring network, n measurements are necessary to recover more than one non-zero element. The key is that every node should be the *endpoint* at least twice, where the endpoints are the nodes at the beginning and the end of a measurement. If node u is an endpoint for at most once, then it is always measured together with one of its neighbors, say v, if ever measured. Then when v is '1', we cannot determine the value of u, either '1' or '0'. Therefore, to recover more than one non-zero element, we need at least $2n$ endpoints, and thus n measurements.

B. Ring with nodes connecting to four closest neighbors

Consider a graph with each node directly connecting to its four closest neighbors as in Fig. 4 (a), denoted by $\mathcal{G}^4$. $\mathcal{G}^4$ is important to the study of small-world networks [35]. $\mathcal{G}^4$ has n more links than the ring network, but we will show that the number of measurements required by compressed sensing to recover k-sparse vectors associated with $\mathcal{G}^4$ significantly reduces from $\Theta(n)$ to $O(k\log(n/k))$. The main idea is referred to as " the use of a hub".

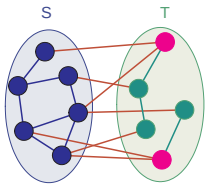


Fig. 3. Hub S for T

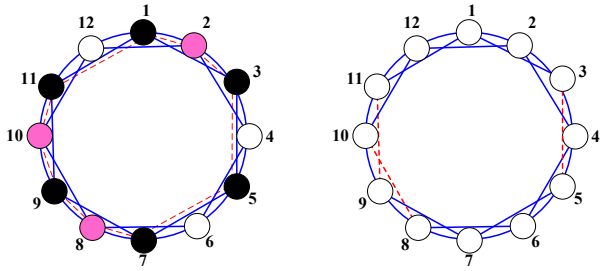
Definition 1. *Given $G=(V,E)$, $S \subseteq V$ is a **hub** for $T \subseteq V$ if G_S is connected, and $\forall u \in T$, $\exists s \in S$ s.t. $(u,s) \in E$.*

We first take one measurement of the sum of nodes in S, denoted by s. For any subset W of T, e.g., the pink nodes in Fig. 3, $S \cup W$ induces a connected subgraph from the hub definition and thus can be measured by one measurement. To measure the sum of nodes in W, we first measure nodes in $S \cup W$ and then subtract s from the sum. Therefore we can apply the measurement constructions for complete graphs on T with this simple modification, and that requires only one additional measurement for the hub S. Thus,

Theorem 2. *With hub S, $M^C_{k,|T|}+1$ measurements are enough to recover k-sparse vectors associated with T.*

The significance of Theorem 2 is that G_T is not necessarily a complete subgraph, i.e., a clique, and it can even be disconnected. As long as there exists a hub S, the measurement construction for a complete graph with the same number of nodes can be applied to T with simple modification. Our later results rely heavily on Theorem 2.

In $\mathcal{G}^4$, if nodes are numbered consecutively around the ring, then the set of all the odd nodes, denoted by T_{o}, form a hub for the set of all the even nodes, denoted by T_{e}. Given a k-sparse vector $\mathbf{x}$, let $\mathbf{x}_{\text{o}}$ and $\mathbf{x}_{\text{e}}$ denote the subvectors of $\mathbf{x}$ with odd and even indices. Then $\mathbf{x}_{\text{o}}$ and $\mathbf{x}_{\text{e}}$ are both at most k-sparse. From Theorem 2, $M^C_{k,\lfloor n/2\rfloor}+1$ measurements are



(a) Measure nodes 2,8 and 10 via hub T_o (b) Delete h long links

Fig. 4. Sparse recovery on graph $\mathcal{G}^4$

enough to recover $\mathbf{x}_e \in \mathcal{R}^{\lfloor n/2 \rfloor}$. Similarly, we can use T_e as a hub to recover the subvector $\mathbf{x}_o \in \mathcal{R}^{\lceil n/2 \rceil}$ with $M^C_{k,\lceil n/2 \rceil} + 1$ measurements, and thus $\mathbf{x}$ is recovered.

Corollary 1. *All k-sparse vectors associated with $\mathcal{G}^4$ can be recovered with $M^C_{k,\lfloor n/2 \rfloor} + M^C_{k,\lceil n/2 \rceil} + 2$ measurements, which is $O(2k \log(n/(2k)))$.*

From ring network to $\mathcal{G}^4$, although the number of links only increases by n, the number of measurements required to recover k-sparse vectors significantly reduces from $\Theta(n)$ to $O(2k \log(n/(2k)))$. This value is in the same order as $M^C_{k,n}$, while the number of links in $\mathcal{G}^4$ is only $2n$, compared with $n(n-1)/2$ links in a complete graph.

Moreover, our estimate $O(2k \log(n/(2k)))$ on the minimum number of measurements required to recover k-sparse vectors greatly improves over the existing results in [13], [38], both of which are based on the mixing time of a random walk. The mixing time $T(n)$ can be roughly interpreted as the minimum length of a random walk on graph G such that its distribution is close to the stationary distribution on G. [38] proves that $O(kT^2(n) \log n)$ measurements can identify k-sparse vectors with overwhelming probability by compressed sensing. [13] needs $O(k^2 T^2(n) \log(n/k))$ measurements to identify k non-zero elements by group testing. In $\mathcal{G}^4$, $T(n)$ should be at least $n/4$. Then both results provide no saving in the number of measurements for $\mathcal{G}^4$ as the mixing time is $\Theta(n)$.

Besides the explicit measurement construction based on the hub idea, we can also recover k-sparse vectors associated with $\mathcal{G}^4$ with $O(\log n)$ random measurements. We need to point out that these random measurements do not depend on the measurement constructions for a complete graph.

Consider an n-step Markov chain $\{X_k, 1 \le k \le n\}$ with $X_1 = 1$. For any $k \le n-1$, if $X_k = 0$, then $X_{k+1} = 1$; if $X_k = 1$, then X_{k+1} can be 0 or 1 with equal probability. Clearly any realization of this Markov chain does not contain two or more consecutive zeros, and thus is a feasible row of the measurement matrix. Moreover,

Theorem 3. *With high probability all k-sparse vectors associated with $\mathcal{G}^4$ can be recovered with $O(g(k) \log n)$ measurements obtained from the above Markov chain, where $g(k)$ is a function of k.*

Proof: See Appendix. ∎

Adding n links in the form $(i, i+2 (\text{mod } n))$ to the ring network greatly reduces the number of measurements needed from $\Theta(n)$ to $O(\log n)$. Then how many links in the form $(i, i+2 (\text{mod } n))$ shall we add to the ring network such that the minimum number of measurements required to recover k-sparse vectors is exactly $\Theta(\log n)$? The answer is $n - \Theta(\log n)$. To see this, let $\mathcal{G}^4_h$ denote the graph obtained by deleting h links in the form $(i, i+2 (\text{mod } n))$ from $\mathcal{G}^4$. For example in Fig. 4 (b), we delete links $(3,5)$, $(8,10)$ and $(9,11)$ in red dashed lines from $\mathcal{G}^4$. Given h, our following results do not depend on the specific choice of links to remove. We have

Theorem 4. *The minimum number of measurements required to recover k-sparse vectors associated with $\mathcal{G}^4_h$ is lower bounded by $\lceil h/2 \rceil$, and upper bounded by $2M^C_{k,\lceil \frac{n}{2} \rceil} + h + 2$.*

Proof: Let D denote the set of nodes such that for every $i \in D$, link $(i-1, i+1)$ is removed from $\mathcal{G}^4$. The proof of the lower bound follows the proof of Theorem 2 in [32]. The key idea is that recovering one non-zero element in D is equivalent to recovering one non-zero element in a ring network with h nodes, and thus $\lceil h/2 \rceil$ measurements are necessary.

For the upper bound, we first measure nodes in D separately with h measurements. Let S contain the even nodes in D and all the odd nodes. S can be used as a hub to recover the k-sparse subvectors associated with the even nodes that are not in D, and the number of measurements used is at most $M^C_{k,\lfloor \frac{n}{2} \rfloor} + 1$. We similarly recover k-sparse subvectors associated with odd nodes that are not in D using the set of the odd nodes in D and all the even nodes as a hub. The number of measurements is at most $M^C_{k,\lceil \frac{n}{2} \rceil} + 1$. Sum them up and the upper bound follows. ∎

Together with (1), Theorem 4 directly implies that if $\Theta(\log n)$ links in the form $(i, i+2 (\text{mod } n))$ are deleted from $\mathcal{G}^4$, then $\Theta(\log n)$ measurements are both necessary and sufficient to recover associated k-sparse vectors for any constant k.

Since the number of measurements required by compressed sensing is greatly reduced when we add n links to the ring network, one may wonder whether the number of measurements needed by group testing can be greatly reduced or not. Our next result shows that this is not the case for group testing, please refer to Appendix for its proof.

Proposition 1. $\lfloor n/4 \rfloor$ *measurements are necessary to locate two non-zero elements associated with $\mathcal{G}^4$ by group testing.*

By Corollary 1 and Proposition 1, we observe that in $\mathcal{G}^4$, with compressed sensing the number of measurements needed to recover k-sparse vectors is $O(2k \log(n/(2k)))$, while with group testing, $\Theta(n)$ measurements are required if $k \ge 2$.

C. Two-dimensional grid

Next we consider the two-dimensional grid, denoted by $\mathcal{G}^{2d}$. $\mathcal{G}^{2d}$ has $\sqrt{n}$ rows and $\sqrt{n}$ columns. We assume $\sqrt{n}$ to be even here, and also skip '$\lceil \cdot \rceil$' and '$\lfloor \cdot \rfloor$' for notational simplicity.

The idea of measurement construction is still the use of a hub. First, Let S_1 contain the nodes in the first row and all the nodes in the odd columns, i.e., the black nodes in Fig. 5. Then S_1 can be used as a hub to measure k-sparse subvectors associated with nodes in $V \backslash S_1$. The number of measurements

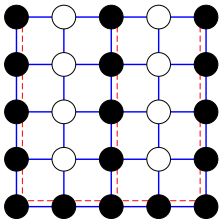


Fig. 5. Two-dimensional grid

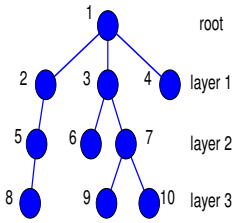


Fig. 6. Tree topology

is $M^C_{k,(n/2-\sqrt{n}/2)}+1$. Then let S_2 contain the nodes in the first row and all the nodes in the even columns, and use S_2 as a hub to recover up to k-sparse subvectors associated with nodes in $V\backslash S_2$. Then number of measurements required is also $M^C_{k,(n/2-\sqrt{n}/2)}+1$. Finally, use nodes in the second row as a hub to recover sparse subvectors associated with nodes in the first row. Since nodes in the second row are already identified in the above two steps, then we do not need to measure the hub separately in this step. The number of measurements here is $M^C_{k,\sqrt{n}}$. Therefore,

With $2M^C_{k,n/2-\sqrt{n}/2}+M^C_{k,\sqrt{n}}+2$ measurements one can ecover k-sparse vectors associated with $\mathcal{G}^{2d}$.

D. Tree

Next we consider a tree topology as in Fig. 6. For a given tree, the root is treated as the only node in layer 0. The nodes that are t steps away from the root are in layer t. We say the tree has depth h if the farthest node is h steps away from the root. Let n_i denote the number of nodes on layer i, and $n_0=1$. We construct measurements to recover vectors associated with a tree by the following *tree approach*.

We recover the nodes layer by layer starting from the root, and recovering nodes in layer i requires that all the nodes above layer i should already be recovered. First measure the root separately. When recovering the subvector associated with nodes in layer i ($2\le i\le h$), we can measure the sum of any subset of nodes in layer i using some nodes in the upper layers as a hub and then delete the value of the hub from the obtained sum. One simple way to find a hub is to trace back from nodes to be measured on the tree simultaneously until they reach one same node. For example in Fig. 6, to measure the sum of nodes 5 and 7, we trace back to the root and measure the sum of nodes 1, 2, 3, 5, and 7 and then subtract the values of nodes 1, 2, and 3, which are already identified when we recover nodes in the upper layers. With this approach, we have,

$\sum_{i=0}^{h} M^C_{k,n_i}$ measurements are enough to recover k-sparse vectors associated with a tree with depth h, where n_i is the number of nodes in layer i.

IV. Sparse Recovery over General Graphs

In this section we consider recovering k-sparse vectors associated with general graphs. The graph is assumed to be connected. If not, we design measurements to recover k-sparse subvectors associated with each component separately.

In Section IV-A we propose a general design guideline based on "r-partition". The key idea is to divide the nodes into a small number of groups such that each group can be measured with the help of a hub. Since finding the minimum number of such groups turns out to be NP-hard in general, in Section IV-B we propose a simple algorithm to design measurements on any given graph.

A. Measurement Construction Based on r-partition

Definition 2 (r-partition)**.** *Given $G=(V,E)$, disjoint subsets N_i ($i=1,...,r$) of V form an r-**partition** of G if and only if these two conditions both hold: (1) $\cup_{i=1}^{r}N_i=V$, and (2) $\forall i$, $V\backslash N_i$ is a hub for N_i.*

Clearly, $T_{\rm o}$ and $T_{\rm e}$ form a 2-partition of graph $\mathcal{G}^4$. With Definition 2 and Theorem 2, we have

Theorem 5. *If G has an r-partition N_i ($i=1,...,r$), then the number of measurements needed to recover k-sparse vectors associated with G is at most $\sum_{i=1}^{r}M^C_{k,|N_i|}+r$, which is $O(rk\log(n/k))$.*

Another example of the existence of an r-partition is the Erdős-Rényi random graph $G(n,p)$ with $p>\log n/n$. The number of our constructed measurements on $G(n,p)$ is less than the existing estimates in [13], [38]. Please refer to Section V for the detailed discussion.

Clearly, if an r-partition exists, the number of measurements also depends on r. In general one wants to reduce r so as to reduce the number of measurements. Given graph G and integer r, the question that whether or not G has an r-partition is called *r-partition problem*. In fact,

Theorem 6. *$\forall r\ge 3$, r-partition problem is NP-complete.*

Please refer to Appendix for its proof. We remark that we cannot prove the hardness of the 2-partition problem though we conjecture it is also a hard problem.

B. Measurement Construction Algorithm for General Graphs

Section IV-A proposes the r-partition concept as a measurement design guideline. But finding an r-partition with the smallest r in general is NP-hard. Given a connected graph G, how shall we efficiently design a small number of measurements to recover k-sparse vectors associated with G?

One simple way is to find the spanning tree of G, and then use the tree approach in Section III-D. The depth of the spanning tree is at least R, where $R=\min_{u\in V}\max_{v\in V}d_{uv}$ is the radius of G with d_{uv} as the length of the shortest path between u and v. This approach only uses links in the spanning tree, and the number of measurements needed is large when the radius R is large. For example, the radius of $\mathcal{G}^4$ is $n/4$, then the tree approach uses at least $n/4$ measurements, while $O(2k\log(n/2k))$ measurements are already enough if we take advantage of the additional links not in the spanning tree.

Here we propose a simple algorithm to design the measurements for general graphs. The algorithm combines the ideas of the tree approach and the r-partition. We still divide nodes into a small number of groups such that each group can be identified via some hub. Here nodes in the same group are the leaf nodes of a spanning tree of a gradually reduced graph. A leaf node has no children on the tree.

Subroutine 1 Leaves(G, u)

Input: graph G, root u

Find a spanning tree T of G rooted at u by breadth-first search, and let S denote the set of leaf nodes of T.

Return: S

Subroutine 2 Reduce(G, u, K)

Input: $G = (V, E)$, H_e for each $e \in E$, and node u

$V = V \backslash u$.
for each two different neighbors v and w of u **do**
  if $(v, w) \notin E$ **then**
    $E = E \cup (v, w)$, $K_{(v,w)} = K_{(v,u)} \cup K_{(u,w)} \cup \{u\}$.
  end if
end for
Return: G, K

Let $G^* = (V^*, E^*)$ denote the orginal graph. The algorithm is built on the following two subroutines. **Leaves**(G, u) returns the set of leaf nodes of a spanning tree of G rooted at u. **Reduce**(G, u, K) deletes u from G and fully connects all the neighbors of u. Specifically, for every two neighbors v and w of u, we add a link (v, w), if not already exist, and let $K_{(v,w)} = K_{(v,u)} \cup K_{(u,w)} \cup \{u\}$, where for each link (s, t), $K_{(s,t)}$ denotes the set of nodes, if any, that connects s and t in the original graph G^*. We record K such that measurements constructed on a reduced graph G can be feasible in G^*.

Given graph G^*, let u denote the node such that $\max_{v \in V^*} d_{uv} = R$, where R is the radius of G^*. Pick u as the root and obtain a spanning tree T of G^* by breadth-first search. Let S denote the set of leaf nodes in T. With $V^* \backslash S$ as a hub, we can design $f(k, |S|) + 1$ measurements to recover up to k-sparse vectors associated with S. We then reduce the network by deleting every node v in S and fully connects its neighbors. For the reduced network G, we repeat the above process until all the nodes are deleted. Note that when designing the measurements in a reduced graph G, if a measurement passes link (v, w), then it should also include nodes in $K_{(v,w)}$ so as to be feasible in the original graph G^*.

In each step tree T is rooted at node u where $\max_{v \in V} d_{uv}$ equals the radius of the current graph G. Since all the leaf nodes of T are deleted in the graph reduction procedure, the radius of the new obtained graph should be reduced by at least one. Then we have at most R iterations in Algorithm 1 until only one node is left. Clearly we have,

Algorithm 1 Measurement construction for graph G^*

Input: $G^* = (V^*, E^*)$.

$G = G^*$, $K_e = \varnothing$ for each $e \in E$
while $|V| > 1$ **do**
  Find the node u such that $\max_{v \in V} d_{uv} = R^G$, where R^G is the radius of G. $S =$**Leaves**(G, u).
  Design $f(k, |S|) + 1$ measurements to recover k-sparse vectors associated with S using nodes in $V \backslash S$ as a hub.
  for each v in S **do**
    $G =$ **Reduce**(G, v, K)
  end for
end while
Measure the last node in V directly.
Output: All the measurements.

Theorem 7. *The number of measurements designed by Algorithm 1 is at most $Rf(k, n) + R + 1$, where R is the radius of the graph.*

We remark that the number of measurements by the spanning tree approach is also no greater than $Rf(k, n) + R + 1$. However, since Algorithm 1 also considers links that are not in the spanning tree, we expect that for general graphs, it uses fewer measurements than the spanning tree approach. This is verified in Experiment 1 in Section VIII.

V. Sparse Recover over Random Graphs

Here we consider measurement constructions over the Erdős-Rényi random graph $G(n, p)$, which has n nodes and every two nodes are connected by a link independently with probability p. The behavior of $G(n, p)$ changes significantly when p varies. We study the dependence of number of measurements needed for sparse recovery on p.

A. $np = \beta \log n$ for some constant $\beta > 1$

Now $G(n, p)$ is connected almost surely [24]. Moreover, we have the following lemma regarding the existence of an r-partition.

Lemma 1. *When $p = \beta \log n/n$ for some constant $\beta > 1$, with probability at least $1 - O(n^{-\alpha})$ for some $\alpha > 0$, every set S of nodes with size $|S| = n/(\beta - \epsilon)$ for any $\epsilon \in (0, \beta - 1)$ forms a hub for the complementary set $T = V \backslash S$, which implies that $G(n, p)$ has a $\lceil \frac{\beta - \epsilon}{\beta - \epsilon - 1} \rceil$-partition.*

Proof: Note that the subgraph G_S is also Erdős-Rényi random graph in $G(n/(\beta - \epsilon), p)$. Since $p = \beta \log n/n > \log(n/(\beta - \epsilon))/(n/(\beta - \epsilon))$, G_S is connected almost surely.

Let P_f denote the probability that there exists some $u \in T$ such that $(u, v) \notin E$ for every $v \in S$. Then

$$\begin{aligned} P_f &= \sum_{u \in T} (1-p)^{|S|} = (1 - \frac{1}{\beta - \epsilon}) n (1 - \beta \log n/n)^{n/(\beta - \epsilon)} \\ &= (1 - 1/(\beta - \epsilon)) n (1 - \beta \log n/n)^{\frac{n}{\beta \log n} \cdot \frac{\beta \log n}{\beta - \epsilon}} \\ &\le (1 - \frac{1}{\beta - \epsilon}) n e^{-\frac{\beta \log n}{\beta - \epsilon}} \quad \le (1 - \frac{1}{\beta - \epsilon}) n^{-\epsilon/(\beta - \epsilon)}. \end{aligned}$$

Thus, S is a hub for T with probability at least $1 - O(n^{-\alpha})$ for $\alpha = \epsilon/(\beta - \epsilon) > 0$. Since the size of T is $(1 - 1/(\beta - \epsilon))n$, $G(n, p)$ has at most $\lceil \frac{\beta - \epsilon}{\beta - \epsilon - 1} \rceil$ such disjoint sets. Then by a simple union bound, one can conclude that $G(n, p)$ has a $\lceil \frac{\beta - \epsilon}{\beta - \epsilon - 1} \rceil$-partition with probability at least $1 - O(n^{-\alpha})$. ■

For example, when $\beta > 2$, Lemma 1 implies that any two disjoint sets N_1 and N_2 with $|N_1| = |N_2| = n/2$ form a 2-partition of $G(n, p)$ with probability $1 - O(n^{-\alpha})$. From Theorem 5 and Lemma 1, and let $\epsilon \to 0$, we have

When $p = \beta \log n/n$ for some constant $\beta > 1$, all k-sparse vectors associated with $G(n, p)$ can be identified with

$O(\lceil \frac{\beta}{\beta-1} \rceil k \log(n/k))$ *measurements with probability at least* $1 - O(n^{-\alpha})$ *for some* $\alpha > 0$.

[13] considers group testing over Erdős-Rényi random graphs and shows that $O(k^2 \log^3 n)$ measurements are enough to identify up to k non-zero entries if it further holds that $p = \Theta(k \log^2 n/n)$. Here with compressed sensing setup and r-partition results, we can recover k-sparse vectors in $\mathcal{R}^n$ with $O(k \log(n/k))$ measurements when $p > \log n/n$. This result also improves over the previous result in [38], which requires $O(k \log^3 n)$ measurements for compressed sensing on $G(n,p)$.

B. $np - \log n \to +\infty$, and $\frac{np-\log n}{\log n} \to 0$

Roughly speaking, p is just large enough to guarantee that $G(n,p)$ is connected almost surely [24]. The diameter $D = \max_{u,v} d_{uv}$ of a connected graph is the greatest distance between any pair of nodes, and here it is concentrated around $\frac{\log n}{\log \log n}$ almost surely [7]. We design measurements on $G(n,p)$ with Algorithm 1. With Theorem 7 and the fact that the radius R is no greater than the diameter D by definition, we have

When $np - \log n \to +\infty$, *and* $\frac{np-\log n}{\log n} \to 0$, $O(k \log n \log(n/k)/\log\log n)$ *measurements can identify* k*-sparse vectors associated with* $G(n,p)$ *almost surely.*

C. $1 < c = np < \log n$

Now $G(n,p)$ is disconnected and has a unique giant component containing $(\alpha + o(1))n$ nodes almost surely with α satisfying $e^{-c\alpha} = 1 - \alpha$, or equivalently,

$$\alpha = 1 - \frac{1}{c} \sum_{k=1}^{\infty} \frac{k^{k-1}}{k!} (ce^{-c})^k,$$

and all the other nodes belong to small components. The expectation of the total number of components in $G(n,p)$ is $(1 - \alpha - c(1-\alpha)^2/2 + o(1))n$ [24]. Since it is necessary to take at least one measurement for each component, $(1-\alpha-c(1-\alpha)^2/2+o(1))n$ is an expected lower bound of measurements required to identify sparse vectors.

The diameter D of a disconnected graph is defined to be the largest distance between any pair of nodes that belong to the same component. Since D is now $\Theta(\log n/\log(np))$ almost surely [14], then for the radius R of the giant component, $R \leq D = O(\log n/\log(np))$, where the second equality holds almost surely. We use Algorithm 1 to design measurements on the giant component, and then measure every node in the small components directly. Thus, k-sparse vectors associated with $G(n,p)$ can be identified almost surely with $O(k \log n \log(n/k)/\log(np)) + (1-\alpha+o(1))n$ measurements.

Note that here almost surely the size of every small component is at most $\frac{\log n + 2\sqrt{\log n}}{np - 1 - \log(np)}$ (Lemma 5, [14]). If $k = \Omega(\log n)$, almost surely $(1 - \alpha + o(1))n$ measurements are necessary to identify subvectors associated with small components, and thus necessary for identifying k-sparse vectors associated with $G(n,p)$. Combing the arguments, we have

When $1 < c = np < \log n$ *with constant* c*, we can identify* k*-sparse vectors associated with* $G(n,p)$ *almost surely with* $O(k \log n \log(n/k)/\log(np)) + (1-\alpha+o(1))n$ *measurements.* $(1-\alpha-c(1-\alpha)^2/2+o(1))n$ *is an expected lower bound of the number of measurements needed. Moreover, if* $k = \Omega(\log n)$, *almost surely* $(1 - \alpha + o(1))n$ *measurements are necessary to identify* k*-sparse vectors.*

D. $np < 1$

Since the expectation of the total number of components in $G(n,p)$ with $np < 1$ is $n - pn^2/2 + O(1)$ [24], then $n - pn^2/2 + O(1)$ is an expected lower bound of the number of measurements required. Since almost surely all components are of size $O(\log n)$, then we need to take n measurements when $k = \Omega(\log n)$. Therefore,

When $np < 1$, *we need at least* $n - pn^2/2 + O(1)$ *measurements to identify* k*-sparse vectors associated with* $G(n,p)$ *in expectation. Moreover, when* $k = \Omega(\log n)$, n *measurements are necessary almost surely.*

VI. Adding additional graph constraints

Our constructions are based on assumptions (A1) and (A2). Here we consider additional graph constraints brought by practical implementation. We first consider measurement construction with length constraint, as measurements with short length are preferred in practice. We then discuss the scenario when only a subset of nodes can act as agents and each measurement should pass at least one agent.

A. Measurements with short length

We have not imposed any constraint on the number of nodes in one measurement. In practice, one may want to take short measurements so as to reduce the communication cost and the measurement noise. We next consider sparse recovery with additional constraint on measurement length, and we discuss two special graphs.

1) Line and Ring: The construction in Section III-A is optimal for a line network in terms of the number of measurements needed, and the length of each measurement is $\lfloor \frac{n+1}{k+1} \rfloor$, which is proportional to n when k is a constant. Here we provide a different construction such that the total number of measurements needed to recover associated k-sparse vectors is $k\lceil \frac{n}{k+1} \rceil + 1$, but each measurement measures at most $k+2$ nodes. We also remark that the number of measurements by this construction is within the minimum plus $\max(k-1, 1)$ for a line network, and the minimum plus k for a ring network.

We construct the measurements as follows. Given k, let B^k be a $k+1$ by $k+1$ square matrix with entries of '1' on the main diagonal and the first row, i.e. $B^k_{ii} = 1$ and $B^k_{1i} = 1$ for all i. If k is even, let $B^k_{i(i-1)} = 1$ for all $2 \leq i \leq k+1$; if k is odd, let $B^k_{i(i-1)} = 1$ for all $2 \leq i \leq k$. $B^k_{ij} = 0$ elsewhere. Let $t = \lceil \frac{n}{k+1} \rceil$, we construct a $(kt+1)$ by $(k+1)t$ matrix A based on B^k. Given set $S \subseteq \{1, ..., kt+1\}$ and set $T \subseteq \{1, ..., (k+1)t\}$, A_{ST} is the submatrix of A with row indices in S and column indices in T. For all $i = 1, ..., t$, let $S_i = \{(i-1)k+1, ..., ik+1\}$, and let $T_i = \{(k+1)(i-1)+1, ..., (k+1)i\}$. Define $A_{S_iT_i} = B^k$ for all i. All the other entries of A are zeros. We keep the first n columns of A as a measurement matrix for the line/ring network with n

nodes. Note that the last one or serval rows of the reduced matrix can be all zeros, and we just delete such rows, and let the resulting matrix be the measurement matrix. For example, when $k = 2$ and $n = 9$, we have $t = 3$, and

$$B^2 = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix},$$

and

$$A = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}. \quad (5)$$

When $k = 3$, and $n = 8$, we have $t = 2$ and

$$B^3 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, A = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Each measurement measures at most $k + 2$ nodes when k is even and at most $k + 1$ nodes when k is odd. We have,

Theorem 8. *The above construction can recover k-sparse vectors associated with a line/ring network with at most $k\lceil \frac{n}{k+1} \rceil + 1$ measurements, which is within the minimum number of measurements needed plus k. And each measurement measures at most $k + 2$ nodes.*

Proof: We only need to prove that all k-sparse vectors in $\mathcal{R}^{(k+1)t}$ can be identified with A, which happens if and only if for every vector $\mathbf{z} \neq \mathbf{0}$ such that $A\mathbf{z} = \mathbf{0}$, $\mathbf{z}$ has at least $2k + 1$ non-zero elements.

If $t = 1$, A a $k+1$ by $k+1$ full rank matrix, and the claim holds trivially. We next consider $t \geq 2$. We prove the case when k is even, and skip the similar proof for odd k.

For each integer t' in $[2, t]$, define a submatrix $A_{t'}$ formed by the first $kt' + 1$ rows and the first $(k + 1)t'$ columns of A. For example, for A in (5), we define

$$A_2 = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}, \quad \text{and } A_3 = A.$$

We will prove by induction on t' that (#) every non-zero vector $\mathbf{z} \in \mathcal{R}^{(k+1)t'}$ such that $A_{t'}\mathbf{z} = \mathbf{0}$ holds has at least $2k + 1$ non-zero elements for every t' in $[2, t]$.

First consider A_2, which is a $(2k + 1) \times (2k + 2)$ matrix. From the last k rows of A_2, one can easily argue that for every $\mathbf{z}$ such that $A_2\mathbf{z} = \mathbf{0}$, its last $k + 1$ entries are either all zeros or all non-zeros. If the last $k + 1$ entries of $\mathbf{z}$ are all zeros, let $\mathbf{z}'$ denote the subvector containing the first $k + 1$ entries of $\mathbf{z}$. Then we have $\mathbf{0} = A_2\mathbf{z} = B^k\mathbf{z}'$. Since B^k is full rank, then $\mathbf{z}' = \mathbf{0}$, which implies that $\mathbf{z} = \mathbf{0}$.

Now consider the case that last $k + 1$ entries of $\mathbf{z}$ are all non-zeros. Since $k + 1$ is odd, the sum of these entries is,

$$\sum_{i=k+2}^{2k+2} z_i = z_{k+2} \neq 0. \quad (6)$$

Let $\mathbf{a}_i^T$ $(i = 1, ..., 2k + 1)$ denote the ith row of A_2. We have

$$\mathbf{a}_{k+1}^T\mathbf{z} = \sum_{i=k}^{2k+2} z_i = 0. \quad (7)$$

Combining (6) and (7), we know that

$$z_k + z_{k+1} = -z_{k+2} \neq 0. \quad (8)$$

Thus, at least one of z_k and z_{k+1} is non-zero. Combining (8) with $\mathbf{a}_1^T\mathbf{z} = 0$, we have one of the first $k - 1$ entries of $\mathbf{z}$ is non-zero. From $\mathbf{a}_i^T\mathbf{z} = 0$ for $2 \leq i \leq k - 1$, one can argue that if one of the first $k - 1$ entries of $\mathbf{z}$ is non-zero, then all the first $k - 1$ entries are non-zero. Therefore, $\mathbf{z}$ has at least $2k + 1$ nonzero entries. (#) holds for A_2.

Now suppose (#) holds for some t' in $[2, t - 1]$. Consider matrix $A_{t'+1}$. Same as the arguments for A_2, one can show that for every $\mathbf{z} \neq \mathbf{0}$ such that $A_{t'+1}\mathbf{z} = \mathbf{0}$, its last $k + 1$ entries are either all zeros or all non-zero. In the former case, let $\mathbf{z}'$ denote the subvector containing the first $(k+1)t'$ entries of $\mathbf{z}$. By induction hypothesis, $\mathbf{z}'$ has at least $2k + 1$ nonzero entries, thus so does $\mathbf{z}$.

If the last $k + 1$ entries of $\mathbf{z}$ are all non-zero, like in the A_2 case, we argue that the sum of $z_{(k+1)t'-1}$ and $z_{(k+1)t'}$ is non-zero, which implies that at least one of them is non-zero. Also consider $\mathbf{a}_i^T\mathbf{z} = 0$ with $i = rk + 1$ for every integer r in $[0, t'-1]$, one can argue that there exist j in $[0, t'-1]$ such that the sum of all $k - 1$ entries from $z_{j(k+1)+1}$ to $z_{j(k+1)+k-1}$ is non-zero. Then, from $\mathbf{a}_i^T\mathbf{z} = 0$ for $i = jk + 2, ..., jk + k - 1$, we know that if the sum of $z_{j(k+1)+1}$ to $z_{j(k+1)+k-1}$ is non-zero, every entry is non-zero. We conclude that in this case $\mathbf{z}$ also has at least $2k + 1$ nonzero entries.

By induction over t', every $\mathbf{z} \neq \mathbf{0}$ such that $A\mathbf{z} = \mathbf{0}$ has at least $2k + 1$ non-zero entries, then the result follows. ■

This construction measures at most $k + 2$ nodes in each measurement. If measurements with constant length are preferred, we provide another construction method such that every measurement only measures at most three nodes. This method requires more measurements, $(2k-1)\lceil \frac{n}{2k} \rceil + 1$ measurements to recover k-sparse vectors associated with a line/ring network.

Given k, let D^k be a $2k$ by $2k$ square matrix having entries of '1' on the main diagonal and the subdiagonal and '0' elsewhere, i.e. $D^k_{ii} = 1$ for all i and $D^k_{i(i-1)} = 1$ for all $i \geq 2$, and $D^k_{ij} = 0$ elsewhere. Let $t = \lceil \frac{n}{2k} \rceil$, we construct a $(2kt - t + 1)$ by $2kt$ matrix A based on D^k. Let $S_i = \{(i - 1)(2k - 1) + 1, ..., i(2k - 1) + 1\}$, and let $T_i = \{2k(i - 1) + 1, ..., 2ki\}$. Define $A_{S_iT_i} = D^k$ for all $i = 1, ..., t$, and $A_{ij} = 0$ elsewhere. We keep the first n columns of A as the measurement matrix. For example, when

$k = 2$ and $n = 8$, we have

$$D^2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix},$$

and

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}. \quad (9)$$

Theorem 9. *The above constructed $(2k-1)\lceil \frac{n}{2k} \rceil + 1$ measurements can identify k-sparse vectors associated with a line/ring network of n nodes, and each measurement measures at most three nodes.*

Proof: When $t = 1$, A is a full rank square matrix. We focus on the case that $t \geq 2$. For each integer t' in $[2, t]$, define a submatrix $A_{t'}$ formed by the first $2kt' - t' + 1$ rows and the first $2kt'$ columns of A. We will prove by induction on t' that every $\mathbf{z} \neq \mathbf{0}$ such that $A_{t'}\mathbf{z} = \mathbf{0}$ holds has at least $2k + 1$ non-zero elements for every t' in $[2, t]$.

First consider A_2. For A in (9), $A_2 = A$. From the first $2k - 1$ rows of A_2, one can check that for every $\mathbf{z}$ such that $A_2\mathbf{z} = \mathbf{0}$, its first $2k - 1$ entries are zeros. From the $2k$th row of A_2, we know that z_{2k} and z_{2k+1} are either both zeros or both non-zero. In the former case, the remaining $2k-1$ entries of $\mathbf{z}$ must be zeros, thus, $\mathbf{z} = \mathbf{0}$. In the latter case, one can check that the remaining $2k - 1$ entries are all non-zero, and therefore $\mathbf{z}$ has $2k + 1$ non-zero entries.

Now suppose the claim holds for some t' in $[2, t - 1]$. Consider vector $\mathbf{z} \neq \mathbf{0}$ such that $A_{t'+1}\mathbf{z} = \mathbf{0}$. If $z_{2kt'+1} = 0$, it is easy to see that the last $2k$ entries of $\mathbf{z}$ are all zeros. Then by induction hypothesis, at least $2k+1$ entries of the first $2kt'$ elements of $\mathbf{z}$ are non-zero. If $z_{2kt'+1} \neq 0$, one can check that the last $2k-1$ entries of $\mathbf{z}$ are all non-zero, and at least one of $z_{2kt'-1}$ and $z_{2kt'}$ is non-zero. Thus, $\mathbf{z}$ also has at least $2k+1$ non-zero entries in this case.

By induction over t', every $\mathbf{z} \neq \mathbf{0}$ such that $A\mathbf{z} = \mathbf{0}$ has at least $2k + 1$ non-zero entries, then the theorem follows. ■

The number of measurements by this construction is greater than those of the previous methods. But the advantage of this construction is that the number of nodes in each measurement is at most three, no matter how large n and k is.

2) Ring with each node connecting to four neighbors: We next consider $\mathcal{G}^4$ in Fig. 4 (a). We further impose the constraint that the number of nodes in each measurement cannot exceed d for some predetermined integer d. We neglect $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ for notational simplicity.

All the even nodes are divided into n/d groups, each of which contains $d/2$ consecutive even nodes. We use each group as a hub to measure $d/2$ odd nodes that have direct links with nodes in the hub. Then we can identify the values related to all the odd nodes with $nM^C_{k,d/2}/d + n/d$ measurements, and the number of nodes in each measurement does not exceed d. We then measure the even nodes with groups of odd nodes as hubs. In total, the number of measurements is $2nM^C_{k,d/2}/d + 2n/d$, which is $O(2kn\log(d/2)/d)$. When d equals to n, the result coincides with Theorem 1. Since n/d measurements are needed to measure all the nodes at least once, we have

Theorem 10. *The number of measurements needed to recover k-sparse vectors associated with $\mathcal{G}^4$ with each measurement containing at most d nodes is lower bounded by n/d, and upper bounded by $O(2kn\log(d/2)/d)$.*

The ratio of the number of measurements by our construction to the minimum number needed with length constraint is within $Ck\log(d/2)$ for some constant C.

B. Measurements passing at least one node in a fixed subset

Recall that for a connected subgraph G_S, agent u in S aggregates the sum of all the nodes in S, and sends the measurement to the central operator. One may want to reduce the number of agents in a network so as to reduce the monitoring and transmission costs. If only a small set Y of nodes can perform as the agents, every constructed measurement is then required to pass at least one node in Y.

This can be achieved with small modification to the construction algorithm. Let Y denote the set of agents. Let D denote the set of nodes that can be measured together via a hub H. If $H \cap Y$ is not empty, every measurement contains at least one node in Y automatically as all the nodes in H are included. Now consider the case that $H \cap Y$ is empty. If there exists a path P from some node j in H to some node f in Y such that $P \cap D$ is empty, then let $\hat{H} := H \cup P$ be the new hub, and design measurements for D using hub $\hat{H}$. Then every measurement contains all nodes in $\hat{H}$ and thus the agent f. If such a path does not exist, pick any node i in D and any node f in Y, find the shortest path P' between i and f. Let $H' := H \cup P'$ be the hub, and let $D' := D\backslash i$ be the set of nodes that can be measured via H'. Then every measurement contains agent f. Node i can be measured by two additional measurements, one measures path P', and the other measures $P'\backslash i$. Therefore, with this simple modification, we can measure the same set of nodes with each measurement contains at least one node in Y, and the total number of measurements increases by at most two.

We summarize the above modification in subroutine **Agent**. For measurement design on general graphs, we first replace step 4 in Algorithm 1 in Section IV-B with subroutine Agent($V\backslash S$, S, Y, G). Then in each iteration the number of measurements is increased by at most two. We then replace step 9 with measuring the paths P^* and $P^*\backslash n_{\text{last}}$, where n_{last} is the last node in G, and P^* connects n_{last} to any node j in Y on the original graph. Therefore, the total number of measurements needed by the modified algorithm is upper bounded by $Rf(k, n) + 3R + 2$, and each measurement in the modified version contains at least one node in Y.

VII. Sensitivity to hub measurement errors

In constructions based on the use of a hub, in order to measure nodes in S using hub H, we first measure the sum

Subroutine 3 Agent(H, D, Y, G)

Input: hub H, set D of nodes, set Y of agents, G

if $H \cap Y \neq \Phi$ **then**
  Design $f(k, |D|)+1$ measurements to recover k-sparse vectors associated with D using H as a hub.
else
  Find the shortest path between every node in H and every node in Y.
  if there exists a shortest path P s.t. $P \cap D = \Phi$ **then**
    Design $f(k, |D|)+1$ measurements to recover nodes in D using $\hat{H} = H \cup P$ as a hub.
  else
    pick a node i in D and a node f in Y, find the shortest path P' between i and f.
    $D' := D \backslash i$, $H' := H \cup P'$, design $f(k, |D'|) + 1$ measurements to recover D' with H' as a hub.
    Measure P' and $P' \backslash i$ to recover node i.
  end if
end if

of nodes in H, and then delete it from other measurements to obtain the sum of some subset of nodes in S. This arises the issue that if the sum of H is not measured correctly, this single error would be introduced into all the measurements. Here we prove that successful recovery is still achievable when a hub measurement is erroneous.

Mathematically, let $\mathbf{x}_S$ denote the sparse vector associated with S, and let $\mathbf{x}_H$ denote the vector associated with H and let $A^{m\times|S|}$ be a measurement matrix that can identify k-sparse vectors associated with a complete graph of $|S|$ nodes. We arrange the vector $\mathbf{x}$ such that $\mathbf{x} = [\mathbf{x}_S^T \quad \mathbf{x}_H^T]^T$, then

$$F = \begin{bmatrix} A & W^{m\times|H|} \\ \mathbf{0}_{|S|}^T & \mathbf{1}_{|H|}^T \end{bmatrix}$$

is the measurement matrix for detecting k non-zeros in S using hub H, where W is a matrix with all '1's, $\mathbf{0}$ is a column vector of all '0's, and $\mathbf{1}$ is a column vector of all '1's. Let vector $\mathbf{z}$ denote the first m measurements, and let z_0 denote the last measurement of the hub H. Then

$$\begin{bmatrix} \mathbf{z} \\ z_0 \end{bmatrix} = \begin{bmatrix} A\mathbf{x}_S + \mathbf{1}^T\mathbf{x}_H\mathbf{1}_m \\ \mathbf{1}^T\mathbf{x}_H \end{bmatrix},$$

or equivalently

$$\mathbf{z} - z_0\mathbf{1}_m = A\mathbf{x}_S. \tag{10}$$

If there is some error e_0 in the last measurement, i.e., instead of z_0, the actual measurement we obtain is

$$\hat{z}_0 = \mathbf{1}^T\mathbf{x}_H + e_0,$$

e_0 hurts the recovery accuracy of $\mathbf{x}_S$ through (10).

To eliminate the impact of e_0, we model it as an entry of an augmented sparse signal to recover. Let $\mathbf{x}' = [\mathbf{x}^T \quad e_0]^T$, and $A' = [A \quad -\mathbf{1}_m]$, we have

$$A'\mathbf{x}' = \mathbf{z} - \hat{z}_0\mathbf{1}_m. \tag{11}$$

Then, recovering $\mathbf{x}_S$ in the presence of hub error e_0 is equivalent to recovering $k+1$-sparse vector $\mathbf{x}'$ from (11).

We consider one special construction of matrix $A^{m\times|S|}$ for a complete graph. A has '1' on every entry in the last row, and takes value '1' and '0' with equal probability independently for every other entry. $A' = [A \quad -\mathbf{1}_m]$, let $\hat{A}$ be the submatrix of the first $m-1$ rows of A'. Let $\mathbf{y} = \mathbf{z} - \hat{z}_0\mathbf{1}_m$, and let $\hat{\mathbf{y}}$ denote the first $m-1$ entries of $\mathbf{y}$. We have,

$$(2\hat{A} - W^{(m-1)\times|S|})\mathbf{x}' = 2\hat{\mathbf{y}} - y_m.$$

We recover $\mathbf{x}'$ by solving the ℓ_1-minimization problem,

$$\min \|\mathbf{x}\|_1, \quad \text{s.t. } (2\hat{A} - W^{(m-1)\times|S|})\mathbf{x} = 2\hat{\mathbf{y}} - y_m. \tag{12}$$

Theorem 11. *With the above construction of A, when $m \geq C(k+1)\log|S|$ for some constant $C > 0$ and $|S|$ is large enough, with probability at least $1 - O(|S|^{-\alpha})$ for some constant $\alpha > 0$, $\mathbf{x}'$ is the unique solution to (12) for all $k+1$-sparse vectors $\mathbf{x}'$ in $\mathcal{R}^{|S|+1}$.*

Theorem 11 indicates that even though the hub measurement is erroneous, one can still identify k-sparse vectors associated with S with $O((k+1)\log|S|)$ measurements.

The proof of Theorem 11 relies heavily on Lemma 2.

Lemma 2. *If matrix $\Phi^{p\times n}$ takes value $-1/\sqrt{p}$ on every entry in the last column and takes value $\pm 1/\sqrt{p}$ with equal probability independently on every other entry, then for any $\delta > 0$, there exists some constant C such that when $p \geq C(k+1)\log n$ and n is large enough, with probability at least $1 - O(n^{-\alpha})$ for some constant $\alpha > 0$ it holds that for every $U \subseteq \{1, ..., n\}$ with $|U| \leq 2k+2$ and for every $\mathbf{x} \in \mathcal{R}^{2k+2}$,*

$$(1-\delta)\|\mathbf{x}\|_2^2 \leq \|\Phi_U\mathbf{x}\|_2^2 \leq (1+\delta)\|\mathbf{x}\|_2^2, \tag{13}$$

where Φ_U is the submatrix of Φ with column indices in U.

Proof: Consider matrix $\Phi'^{p\times n}$ with each entry taking value $\pm 1/\sqrt{p}$ with equal probability independently. For every realization of matrix Φ', construct a matrix $\hat{\Phi}$ as follows. For every $i \in \{1, ..., p\}$ such that $\Phi'_{in} = 1/\sqrt{p}$, let $\hat{\Phi}_{ij} = -\Phi'_{ij}$ for all $j = 1, ..., n$. Let $\hat{\Phi}_{ij} = \Phi'_{ij}$ for every other entry. One can check that $\hat{\Phi}$ and Φ follow the same probability distribution. Besides, according to the construction of $\hat{\Phi}$, for any subset $U \subseteq \{1, ..., n\}$,

$${\Phi'_U}^T\Phi'_U = \hat{\Phi}_U^T\hat{\Phi}_U. \tag{14}$$

The Restricted Isometry Property [10] indicates that the statement in Lemma 2 holds for Φ'. From (14), and the fact that $\|\Phi'_U\mathbf{x}\|_2^2 = \mathbf{x}^T{\Phi'_U}^T\Phi'_U\mathbf{x}$, the statement also holds for $\hat{\Phi}$. Since $\hat{\Phi}$ and Φ follow the same probability distribution, the lemma follows. ■

Proof: (of Theorem 11) From Lemma 2, when $m \geq C(k+1)\log|S|$ for some $C > 0$ and $|S|$ is large enough, with probability at least $1 - O(|S|^{-\alpha})$, matrix $(2\hat{A} - W^{(m-1)\times|S|})/\sqrt{m-1}$ satisfies (13) for some small enough δ, say $\delta < \sqrt{2} - 1$. Then from [11], [26], (12) can recover all $k+1$-sparse vectors correctly. ■

VIII. Simulation

Experiment 1 (Effectiveness of Algorithm 1): Given a graph G, we apply Algorithm 1 to divide the nodes into groups such that each group (except the last one) can be measured

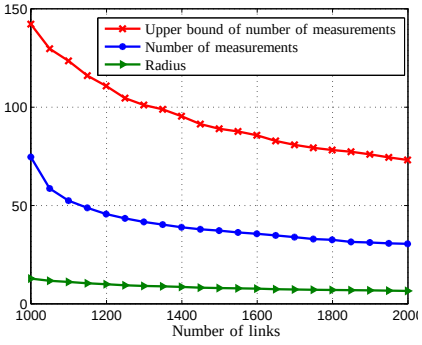


Fig. 6. Random graph with $n = 1000$

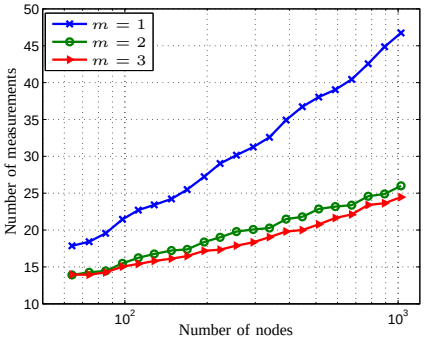


Fig. 7. BA model with increasing n and different m

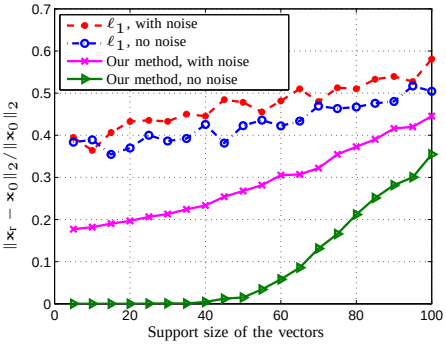


Fig. 8. Recovery performance with hub errors

via some hub. The last group contains one node and can be measured directly. It is know that $M^C_{1,n} = \lceil \log(n+1) \rceil$, and the corresponding measurement matrix has the binary expansion of integer i as column i [21]. Also from (1) the number of measurements required to recovery k-sparse vectors is within a constant times $kM^C_{1,n}$. Therefore, here we design measurements to recover 1-sparse vectors on G as an example. The total number of constructed measurements is $\sum_i^{q-1} \lceil \log(n_i + 1) \rceil + q$, where n_i is the number of nodes in group i and q is the total number of groups.

In Fig. 6, we gradually increase the number of links in a graph with $n = 1000$ nodes. We start with a uniformly generated random tree, and in each step randomly add 25 links to the graph. All the results are averaged over one hundred realizations. The number of measurements constructed decreases from 73 to 30 when the number of links increases from $n-1$ to $2n-1$. Note that the number of measurements is already within $3M^C_{1,n}$ when the average node degree is close to 4. The radius R of the graph decreases from 13 to 7, and we also plot the upper bound $R\lceil \log n \rceil + R + 1$ provided by Theorem 7. One can see that the number of measurements actually constructed is much less than the upper bound.

In Fig. 7, we consider the scale-free network with Barabási-Albert (BA) model [4] where the graph initially has m_0 connected nodes, and each new node connects to m existing nodes with a probability that is proportional to the degree of the existing nodes. We start with a random tree of 10 nodes and increase the total number of nodes from 64 to 1024. Every result is averaged over one hundred realizations. Since the diameter of BA model is $O(\log n / \log\log n))$ [8], then by Theorem 7, the number of our constructed measurements is upper bounded by $O(\log^2 n / \log\log n))$. As the mixing time of BA model is $O(\log n)$ [30], methods in [13] and [38] require $O(\log^3 n)$ random measurements.

Experiment 2 (Recovery Performance with Hub Error): We generate a graph with $n = 500$ nodes from BA model. Algorithm 1 divides nodes into four groups with 375, 122, 2 and 1 node respectively. For each of the first two groups with size n_i $(i = 1, 2)$, we generate $\lceil n_i/2 \rceil$ random measurements each measuring a random subset of the group together with its hub. Every node of the group is included in the random subset independently with probability 0.5. We also measure the two hubs directly. Each of the three nodes in the next two groups is measured directly by one measurement. The generated matrix A is 254 by 500. We generate a sparse vector $\mathbf{x}_0$ with i.i.d. zero-mean Gaussian entries on a randomly chosen support, and normalize $\|\mathbf{x}_0\|_2$ to 1.

To recover $\mathbf{x}_0$ from $\mathbf{y} = A\mathbf{x}_0$, one can run the widely used ℓ_1-minimization [11] to recover the subvectors associated with the first two groups, and the last three entries of $\mathbf{x}_0$ can be obtained from measurements directly. However, as discussed in Section VII, an error in a hub measurement degrades the recovery accuracy of subvectors associated with that group. To address this issue, we use a modified ℓ_1-minimization in which the errors in the two hubs are treated as entries of an augmented vector to recover. Specifically, let the augmented vector $\mathbf{z} = [\mathbf{x}_0^T, e_1, e_2]^T$ and the augmented matrix $\tilde{A} = [A\ \boldsymbol{\beta}\ \boldsymbol{\gamma}]$, where e_1 (or e_2) denotes the error in the measurement of the first (second) hub, and the column vector $\boldsymbol{\beta}$ (or $\boldsymbol{\gamma}$) has '1' in the row corresponding to the measurement of the first (or second) hub and '0' elsewhere. We then recover $\mathbf{z}$ (and thus $\mathbf{x}_0$) from $\mathbf{y} = \tilde{A}\mathbf{z}$ by running ℓ_1-minimization on each group separately.

Fig. 8 compares the recovery performance of our modified ℓ_1-minimization and the conventional ℓ_1-minimization, where the hub errors e_1 and e_2 are drawn from standard Gaussian distribution with zero mean and unit variance. For every support size k, we randomly generate two hundred k-sparse vectors $\mathbf{x}_0$, and let $\mathbf{x}_r$ denote the recovered vector. Even with the hub errors, the average $\|\mathbf{x}_r - \mathbf{x}_0\|_2 / \|\mathbf{x}_0\|_2$ is within 10^{-6} when $\mathbf{x}_0$ is at most 35-sparse by our method, while by ℓ_1-minimization, the value is at least 0.35. We also consider the case that besides errors in hub measurements, every other measurement has i.i.d. Gaussian noise with zero mean and variance 0.04^2. The average $\|\mathbf{x}_r - \mathbf{x}_0\|_2 / \|\mathbf{x}_0\|_2$ here is smaller with our method than that with ℓ_1-minimization.

IX. Conclusion

This paper addresses the sparse recovery problem with graph constraints. We provide explicit measurement constructions for special graphs, and propose measurement design algorithms for general graphs. Our construction for a line network is optimal in terms of the number of measurements needed. The constructions on other graphs also improve over the existing results. We characterize the relationship between the number of measurements for sparse recovery and the

graph topology. We also derive upper and lower bounds of the minimum number of measurements needed for sparse recovery on a given graph. It would be interesting to tighten such bounds, especially the lower bounds.

We have not considered the effect of the measurement noise. Also, we assume full knowledge of the fixed network topology, and measurement construction when the topology is time-varying or partially known is an open question.

APPENDIX

A. Proof of Theorem 3

Let $A^{m\times n}$ denote the matrix with m realizations of the n-step Markov chain. To prove the statement, from [10], we only need to show that the probability that every $2k$ columns of A are linearly independent goes to 1 as n goes to infinity.

Let A_I be a submatrix of A with columns in I, where I is an index set with $|I| = 2k$. Let A_{S_jI} $(1 \le j \le \lfloor \frac{m}{2k} \rfloor)$ be a submatrix of A_I formed by row $2k(j-1)+1$ to row $2kj$ of A_I. Given I, the probability that rank(A_{S_jI})$< 2k$ is the same for every given j, and let it denoted by π_d^I. Let P_d^I denote the probability that rank(A_I)$< 2k$, then

$$P_d^I \le (\pi_d^I)^{\lfloor \frac{m}{2k} \rfloor}. \tag{15}$$

To characterize π_d^I, consider matrix $B^{2k\times 2k}$ with $B_{ii} = 0$ for $i = 2, 3, ..., 2k$ and $B_{ij} = 1$ for all the other elements. Since rank(B)$= 2k$, then

$$\pi_d^I \le 1 - P(A_{S_jI} \text{ is a row permutation of } B). \tag{16}$$

One can check that in this Markov chain, for every $1 \le i < k \le n$, $P(X_k = 1 \mid X_i = 1) \ge 1/2$, $P(X_k = 0 \mid X_i = 1) \ge 1/4$, $P(X_k = 1 \mid X_i = 0) \ge 1/2$, and $P(X_k = 1) \ge 1/2$. Since B has $(2k)!$ different row permutations,

$$P(A_{S_jI} \text{ is a row permutation of } B) \ge (2k)!/2^{4k^2+2k-1}. \tag{17}$$

Combining (15), (16) and (17), we have

$$\begin{aligned} &P(\text{every } 2k \text{ columns of } A \text{ are linearly independent}) \\ =&1 - P(\text{rank}(A_I) < 2k \text{ for some } I \text{ with } |I| = 2k) \\ \ge&1 - \binom{n}{2k} P_d^I \ge 1 - \binom{n}{2k} e^{-(2k)!(\frac{1}{2})^{4k^2+2k-1}\lfloor \frac{m}{2k} \rfloor}, \end{aligned}$$

where the first inequality follows from the union bound. Then if $m = g(k)\log n = (2k+1)2^{4k^2+2k-1}\log n/(2k-1)!$, the probability that every $2k$ columns of A are linearly independent is at least $1 - 1/((2k)!n)$.

B. Proof of Proposition 1

We view nodes $2i-1$ and $2i$ as a group for every i $(1 \le i \le \lfloor \frac{n}{2} \rfloor)$, denoted by B_i. Consider the special case that for some t, both nodes in B_t are '1's, and all other nodes are '0's. Then every measurement that passes either node or both nodes in B_t is always '1'. Consider a reduced graph with B_i, $\forall i$ as nodes, and link (B_i, B_j) $(i \ne j)$ exists only if in $\mathcal{G}^4$ there is a path from a node in B_i to a node in B_j without going though any other node not in B_i or B_j. B_i is '1' if both node $2i-1$ and node $2i$ in $\mathcal{G}^4$ are '1's and is '0' otherwise. The reduced network is a ring with $\lfloor \frac{n}{2} \rfloor$ nodes, and thus $\lfloor n/4 \rfloor$ measurements are required to locate one non-zero element in the reduced network. Then only to locate two consecutive non-zero elements associated with $\mathcal{G}^4$, we need at least $\lfloor n/4 \rfloor$ measurements, and the claim follows.

C. Proof of Theorem 6

Since checking whether or not r given sets form an r-partition takes polynomial time, r-partition problem is NP.

We next show that the NP-complete r-coloring ($r \ge 3$) problem is polynomial time reducible to r-partition problem.

Let $G = (V, E)$ and an integer r be an instance of r-coloring. For every $(u, v) \in E$, add a node w and two links (w, u) and (w, v). Let W denote the set of nodes added. Add a link between every pair of nodes in V not already joined by a link. Let H denote the augmented graph and let $V' = V \cup W$. We claim that if there exists an r-partition of H, then we can obtain an r-coloring of G, and vice versa.

Let S_i $(i = 1, ..., r)$ be an r-partition of H. Suppose there exists link $(u, v) \in E$ s.t. u and v both belong to S_i for some i. Let w denote the node in W that only directly connects to u and v. If $w \in S_i$, then w has both neighbors in the same set with w, contradicting the definition of r-partition. If $w \notin S_i$, then $H_{V'\backslash S_i}$ is disconnected since w does not connect to any node in $V'\backslash S_i$. It also contradicts the definition of r-partition. Thus, for every $(u, v) \in E$, node u and v belong to two sets S_i and S_j with $i \ne j$. Then we obtain an r-coloring of G.

Let $C_i \subset V$ $(i = 1, ..., r)$ denote an r-coloring of G. We claim that $N_i = C_i$ $(i = 1, ..., r-1)$, and $N_r = C_r \cup W$ form an r-partition of H. First note for every $u \in V$, at least one of its neighbors is not in the same set as u. For every $w \in W$, w is directly connected to u and v for some $(u, v) \in E$, and u and v are in different sets C_i and C_j for some $i \ne j$. Therefore, w has at least one neighbor that is not in N_r. Second, we will show $H_{V'\backslash N_i}$ is connected for all i. $H_{V'\backslash N_r}$ is a complete subgraph, and thus connected. For every $i < r$, let $S_i := V\backslash C_i$, then $V'\backslash N_i = S_i \cup W$. H_{S_i} is a complete subgraph, and thus connected. For every $w \in W$, since its two neighbors cannot be both in C_i, then at least one neighbor belongs to S_i, thus $H_{V'\backslash N_r} = H_{S_i \cup W}$ is connected. N_i $(i = 1, ..., r)$ thus forms an r-partition.

REFERENCES

[1] S. S. Ahuja, S. Ramasubramanian, and M. M. Krunz, "Single-link failure detection in all-optical networks using monitoring cycles and paths," *IEEE/ACM Trans. Netw.*, vol. 17, no. 4, pp. 1080–1093, 2009.

[2] L. Applebaum, S. D. Howard, S. Searle, and R. Calderbank, "Chirp sensing codes: Deterministic compressed sensing measurements for fast recovery," *Applied and Computational Harmonic Analysis*, vol. 26, no. 2, pp. 283 – 290, 2009.

[3] P. Babarczi, J. Tapolcai, and P.-H. Ho, "Adjacent link failure localization with monitoring trails in all-optical mesh networks," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 907 –920, 2011.

[4] A. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.

[5] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss., "Combining geometry and combinatorics: a unified approach to sparse signal recovery," *arxiv:0804.4666*, 2008.

[6] T. Blumensath, "Compressed sensing with nonlinear observations," Tech. Rep., 2010.

[7] B. Bollobas, *Random Graphs*, 2nd ed. Cambridge Univ. Press, 2001.

[8] B. Bollobás and O. Riordan, "The diameter of a scale-free random graph," *Combinatorica*, vol. 24, pp. 5–34, 2004.
[9] T. Bu, N. Duffield, F. L. Presti, and D. Towsley, "Network tomography on general topologies," in *Proc ACM SIGMETRICS*, 2002, pp. 21–30.
[10] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
[11] ——, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
[12] Y. Chen, D. Bindel, H. H. Song, and R. Katz, "Algebra-based scalable overlay network monitoring: Algorithms, evaluation, and applications," *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1084 –1097, 2007.
[13] M. Cheraghchi, A. Karbasi, S. Mohajer, and V. Saligrama, "Graph-constrained group testing," *IEEE Trans. Inf. Theory*, vol. 58, no. 1, pp. 248–262, Jan. 2012.
[14] F. Chung and L. Lu, "The diameter of sparse random graphs," *Advances in Applied Mathematics*, vol. 26, no. 4, pp. 257 – 279, 2001.
[15] A. Coates, A. Hero III, R. Nowak, and B. Yu, "Internet tomography," *IEEE Signal Processing Magazine*, vol. 19, no. 3, pp. 47 –65, 2002.
[16] M. Coates, Y. Pointurier, and M. Rabbat, "Compressed network monitoring for ip and all-optical networks," in *Proc. ACM SIGCOMM IMC*, 2007, pp. 241–252.
[17] G. Cormode and S. Muthukrishnan, "Combinatorial algorithms for compressed sensing," ser. Lecture Notes in Computer Science, 2006, vol. 4056, pp. 280–294.
[18] R. DeVore, "Deterministic constructions of compressed sensing matrices," *Journal of Complexity*, vol. 23, no. 4-6, pp. 918 – 925, 2007.
[19] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
[20] D. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," in *Proc. Natl. Acad. Sci. U.S.A.*, vol. 102, no. 27, 2005, pp. 9446–9451.
[21] R. Dorfman, "The detection of defective members of large populations," *Ann. Math. Statist.*, vol. 14, pp. 436–440, 1943.
[22] D.-Z. Du and F. K. Hwang, *Combinatorial Group Testing and Its Applications (Applied Mathematics)*, 2nd ed. World Scientific Publishing Company, 2000.
[23] N. Duffield, "Network tomography of binary network performance characteristics," *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5373 – 5388, 2006.
[24] P. Erdős and A. Rényi, "On the evolution of random graphs," *Publ. Math. Inst. Hung. Acad. Sci.*, pp. 17–61, 1960.
[25] M. Firooz and S. Roy, "Link delay estimation via expander graphs," *arxiv:1106.0941*, 2011.
[26] S. Foucart and M.-J. Lai, "Sparsest solutions of underdetermined linear systems via ℓ_q-minimization for $0 < q \leq 1$," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 395 – 407, 2009.
[27] A. Gopalan and S. Ramasubramanian, "On identifying additive link metrics using linearly independent cycles and paths," *To appear in IEEE/ACM Trans. Netw.*
[28] N. Harvey, M. Patrascu, Y. Wen, S. Yekhanin, and V. Chan, "Non-adaptive fault diagnosis for all-optical networks via combinatorial group testing on graphs," in *Proc. IEEE INFOCOM*, 2007, pp. 697 –705.
[29] J. Haupt, W. Bajwa, M. Rabbat, and R. Nowak, "Compressed sensing for networked data," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 92 –101, 2008.
[30] M. Mihail, C. Papadimitriou, and A. Saberi, "On certain connectivity properties of the internet topology," *Journal of Computer and System Sciences*, vol. 72, no. 2, pp. 239–251, 2006.
[31] H. X. Nguyen and P. Thiran, "Using end-to-end data to infer lossy links in sensor networks," in *Proc. IEEE INFOCOM*, 2006, pp. 1 –12.
[32] J. Tapolcai, B. Wu, P.-H. Ho, and L. Rónyai, "A novel approach for failure localization in all-optical mesh networks," *IEEE/ACM Trans. Netw.*, vol. 19, pp. 275–285, 2011.
[33] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, H. Mobahi, and Y. Ma, "Towards a practical face recognition system: Robust alignment and illumination by sparse representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, no. 99, pp. 1–14, 2011.
[34] M. Wang, W. Xu, E. Mallada, and A. Tang, "Sparse recovery with graph constraints: Fundamental limits and measurement construction," in *Proc. IEEE INFOCOM*, 2012, pp. 1871 –1879.
[35] D. Watts and S. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, 1998.
[36] B. Wu, P.-H. Ho, J. Tapolcai, and X. Jiang, "A novel framework of fast and unambiguous link failure localization via monitoring trails," in *Proc. IEEE INFOCOM*, 2010, pp. 1 –5.
[37] W. Xu and B. Hassibi, "Efficient compressive sensing with deterministic guarantees using expander graphs," in *Proc. IEEE ITW*, 2007, pp. 414 –419.
[38] W. Xu, E. Mallada, and A. Tang, "Compressive sensing over graphs," in *Proc. IEEE INFOCOM*, 2011.
[39] Y. Zhao, Y. Chen, and D. Bindel, "Towards unbiased end-to-end network diagnosis," in *Proc. SIGCOMM*, 2006, pp. 219–230.



Meng Wang (S'06) received the B.E. (Hon.) and M.S. (Hon.) in Electrical Engineering from Tsinghua University, Beijing, China, in 2005 and 2007.

She is currently working towards the Ph.D degree in the School of Electrical and Computer Engineering at Cornell University. Her research interests include communication networks, signal processing, and nonconvex optimization and its applications.

Ms. Wang is a recipient of Jacobs Fellowship of Cornell University in 2008 and 2010.



Weiyu Xu received his B.E. in Information Engineering from Beijing University of Posts and Telecommunications in 2002, and a M.S. degree in Electronic Engineering from Tsinghua University in 2005. He received a M.S. and a Ph.D. degree in Electrical Engineering in 2006 and 2009 from California Institute of Technology, with a minor in Applied and Computational Mathematics.

He is currently an Assistant Professor at the Department of Electrical and Computer Engineering at the University of Iowa. His research interests are in signal processing, compressive sensing, communication networks, information and coding theory.

Dr. Xu is a recipient of the Information Science and Technology Fellowship at Caltech, and the recipient of Charles and Ellen Wilts doctoral research award in 2010.



Enrique Mallada (S'09) received his Ingeniero en Telecomunicaciones degree from Universidad ORT, Uruguay, in 2005. From 2006 to 2008 he was teaching and research assistant in the Department of Telecommunications at Universidad ORT. He is currently pursuing the Ph.D. degree in electrical and computer engineering at Cornell University, Ithaca, NY. His research interests include communication networks, control, nonlinear dynamics and optimization.



Ao Tang (S'01-M'07-SM'11) received the B.E. (Honors.) in electronics engineering from Tsinghua University, Beijing, China, and the M.S. and Ph.D. degrees in electrical engineering with a minor in applied and computational mathematics from the California Institute of Technology, Pasadena, CA, in 1999, 2002, and 2006, respectively.

He is currently an Assistant Professor in the School of Electrical and Computer Engineering at Cornell University, where he conducts research on the control and optimization of engineering networks including communication networks, power networks and on-chip networks.

Dr. Tang was a recipient of the Cornell Engineering School Michael Tien' 72 Excellence in Teaching Award in 2011, and the Young Investigator Award from the Airforce Office of Scientific Research (AFOSR) in 2012.